40° South Guard

Guard alongside Cloudflare AI Gateway

Different problems, different products

Version 1.0 | May 2026

Why this comes up

Cloudflare AI Gateway is a solid product for content moderation, caching, rate limiting, and observability across AI providers. It sits in the request path and gives engineering teams visibility into how their AI integrations behave.

The question we hear is: "Cloudflare already filters content and logs calls. Why do we need Guard?"

The short answer: content moderation is not compliance. Cloudflare answers "is this output safe?" Guard answers "can we prove to a regulator that our controls are active on every call?" They are different questions with different evidence requirements.

Side by side

Capability	40 South Guard	Cloudflare AI Gateway
Primary purpose	Per-call compliance evidence and attestation for all AI providers	Content moderation, caching, rate limiting, and AI observability
Architecture	Inline API proxy with compliance engine, attestation, and evidence vault	Inline proxy with content filtering, logging, and caching
Australian PII	Purpose-built detection with checksum validation for TFN, Medicare, ABN, and ACN	No Australian-specific PII types. General content filtering only.
Regulatory mapping	Every call mapped to CPS 234, the AI Safety Standard, and the Privacy Act	No regulatory mapping. Content categories only (toxicity, violence, etc.).
Per-call attestation	Cryptographically signed attestation per API call with tamper-evident 7-year audit trail	Logging and analytics. No signed attestation or tamper-evident evidence.

AI-specific controls	Prompt injection detection, model inventory, shadow AI blocking, policy enforcement	Content filtering, rate limiting, request caching, and usage analytics
Data residency	Australian-hosted infrastructure. Data never leaves AU jurisdiction.	Global CDN. Data routed through nearest Cloudflare edge node. Customer configures region hints.

Could you run them together?

Yes. Cloudflare AI Gateway is a strong choice for content safety, caching, and operational observability. Guard is the compliance layer that sits alongside it.

A practical deployment would use Cloudflare for content moderation and rate limiting, and Guard for regulatory evidence, Australian PII detection, and per-call attestation. Together, you get both operational safety and audit-ready compliance.

Suggested next step

We can run a 30-minute proof of value showing Guard generating compliance evidence alongside your Cloudflare AI Gateway deployment. You will see per-call attestation and Australian PII detection working on the same API calls that Cloudflare is already moderating.

Contact

hello@40south.au

40south.au

Sources

1. APRA Prudential Standard CPS 234 Information Security, July 2019. www.apra.gov.au/sites/default/files/cps_234_july_2019_for_public_release.pdf
2. Cloudflare, "AI Gateway documentation." developers.cloudflare.com/ai-gateway/
3. Australian Government, "Voluntary AI Safety Standard," August 2024. www.industry.gov.au/publications/voluntary-ai-safety-standard
4. OAIC, "Chapter 8: APP 8 — Cross-border disclosure of personal information." www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-8-app-8-cross-border-disclosure-of-personal-information
5. 40 South Guard technical documentation. 40south.au